



oVirt Gluster Integration

FOSDEM 2013

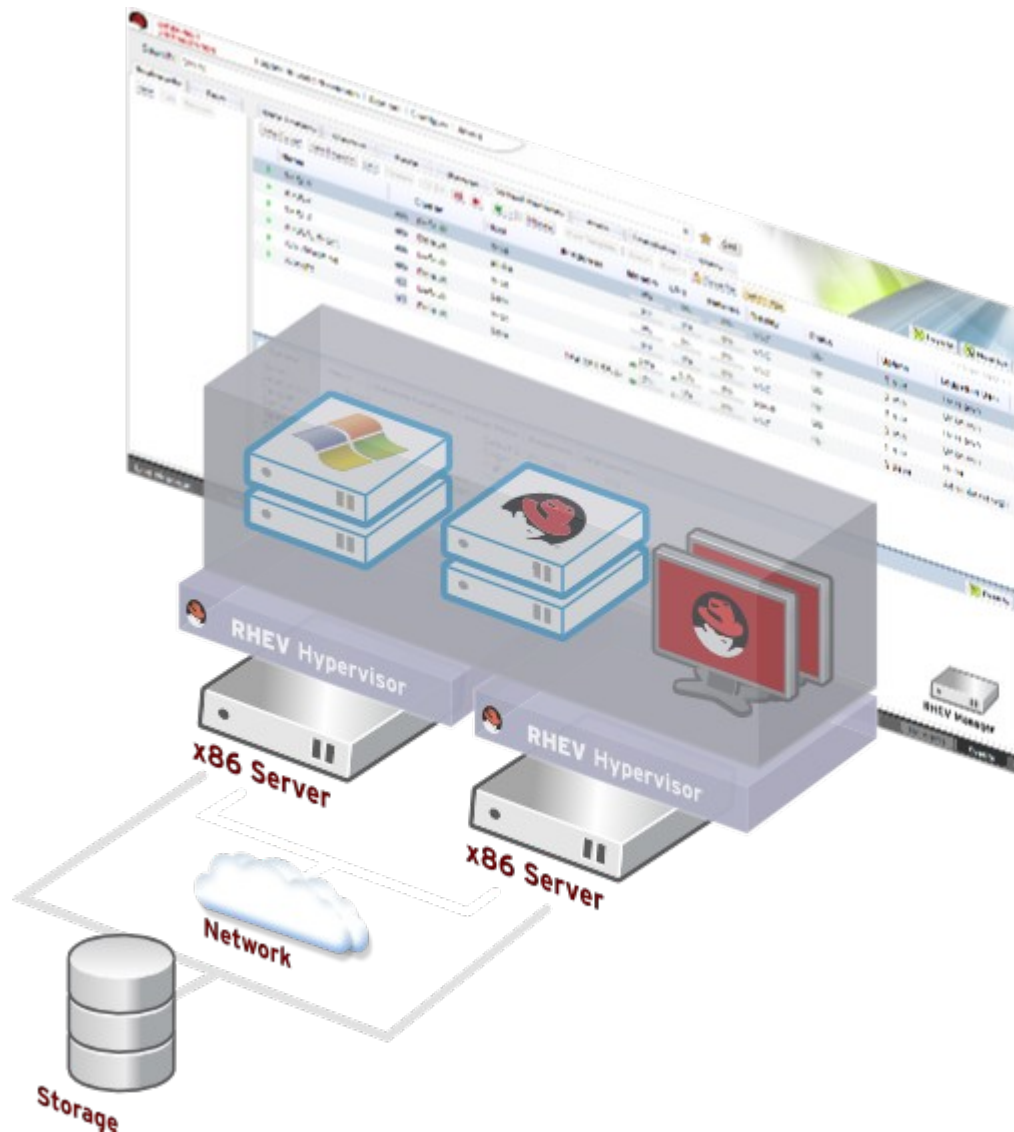
Doron Fediuck
Red Hat



Agenda


- About oVirt and Gluster
- oVirt Gluster Integration
 - Architecture
 - UI
 - REST API
 - CLI
- Coming next

oVirt



- Large scale, centralized management for server and desktop virtualization
- Open source alternative to vCenter / vSphere
- Focus on KVM
- www.ovirt.org
- #ovirt on OFTC

Gluster

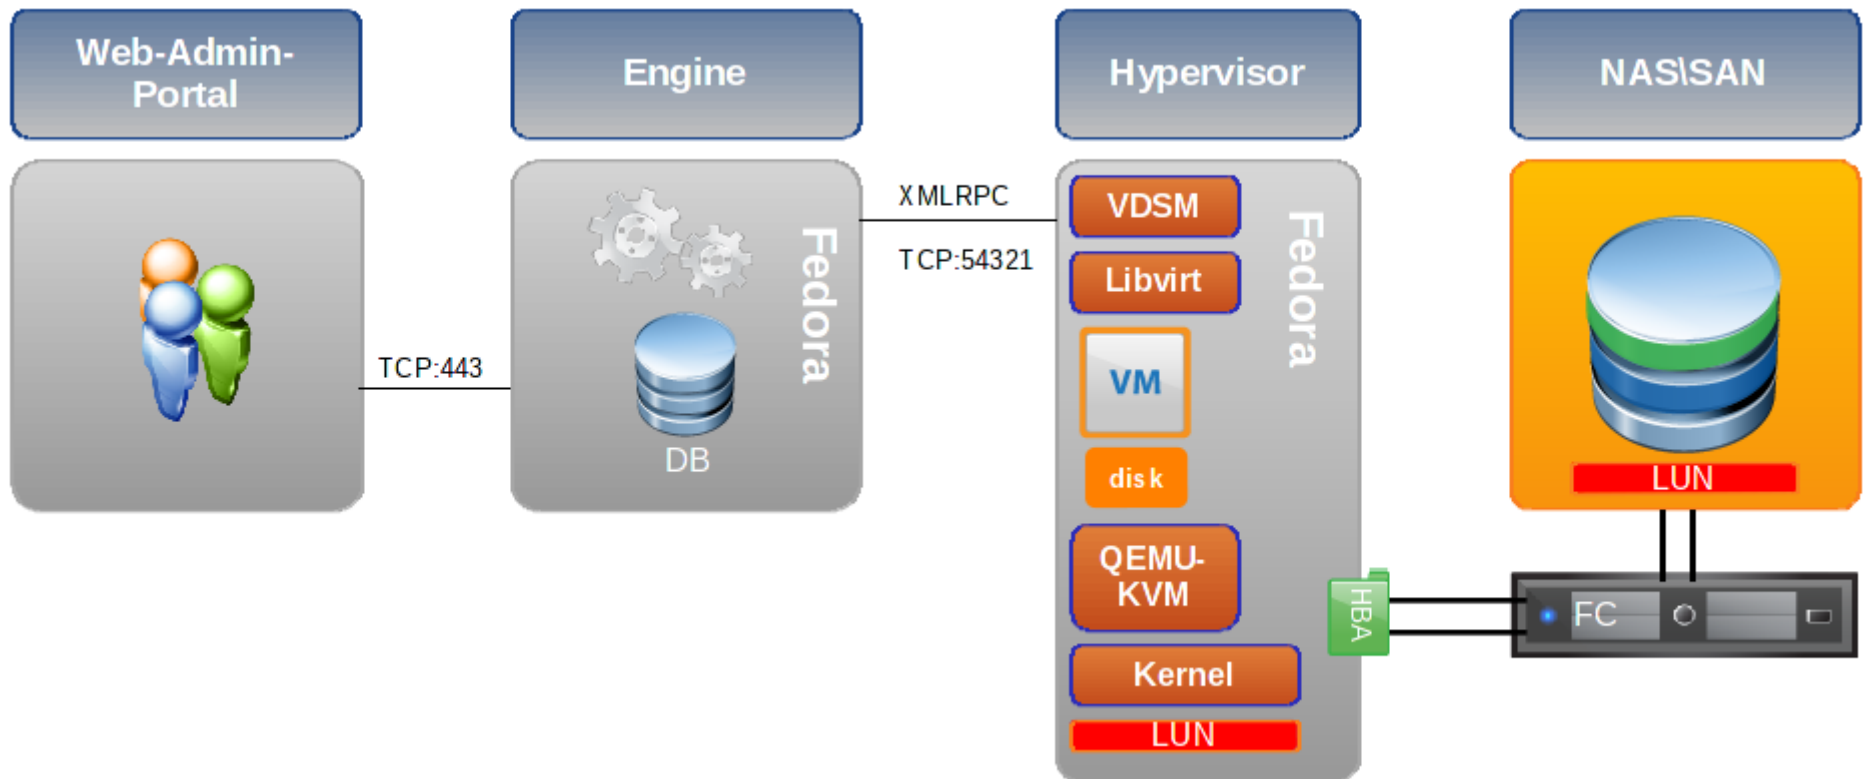


- User-space file system
- Global namespace
- Scale-out clustered storage building blocks
- Supports thousands of clients
- Access using various protocols
- Linear performance scaling
- www.gluster.org
- #gluster on freenode




Architectures

oVirt basic architecture



Overview of oVirt storage concepts



GlusterFS Concepts



VOLUME

is a namespace presented as a POSIX mount point and is comprised of bricks.



BRICK

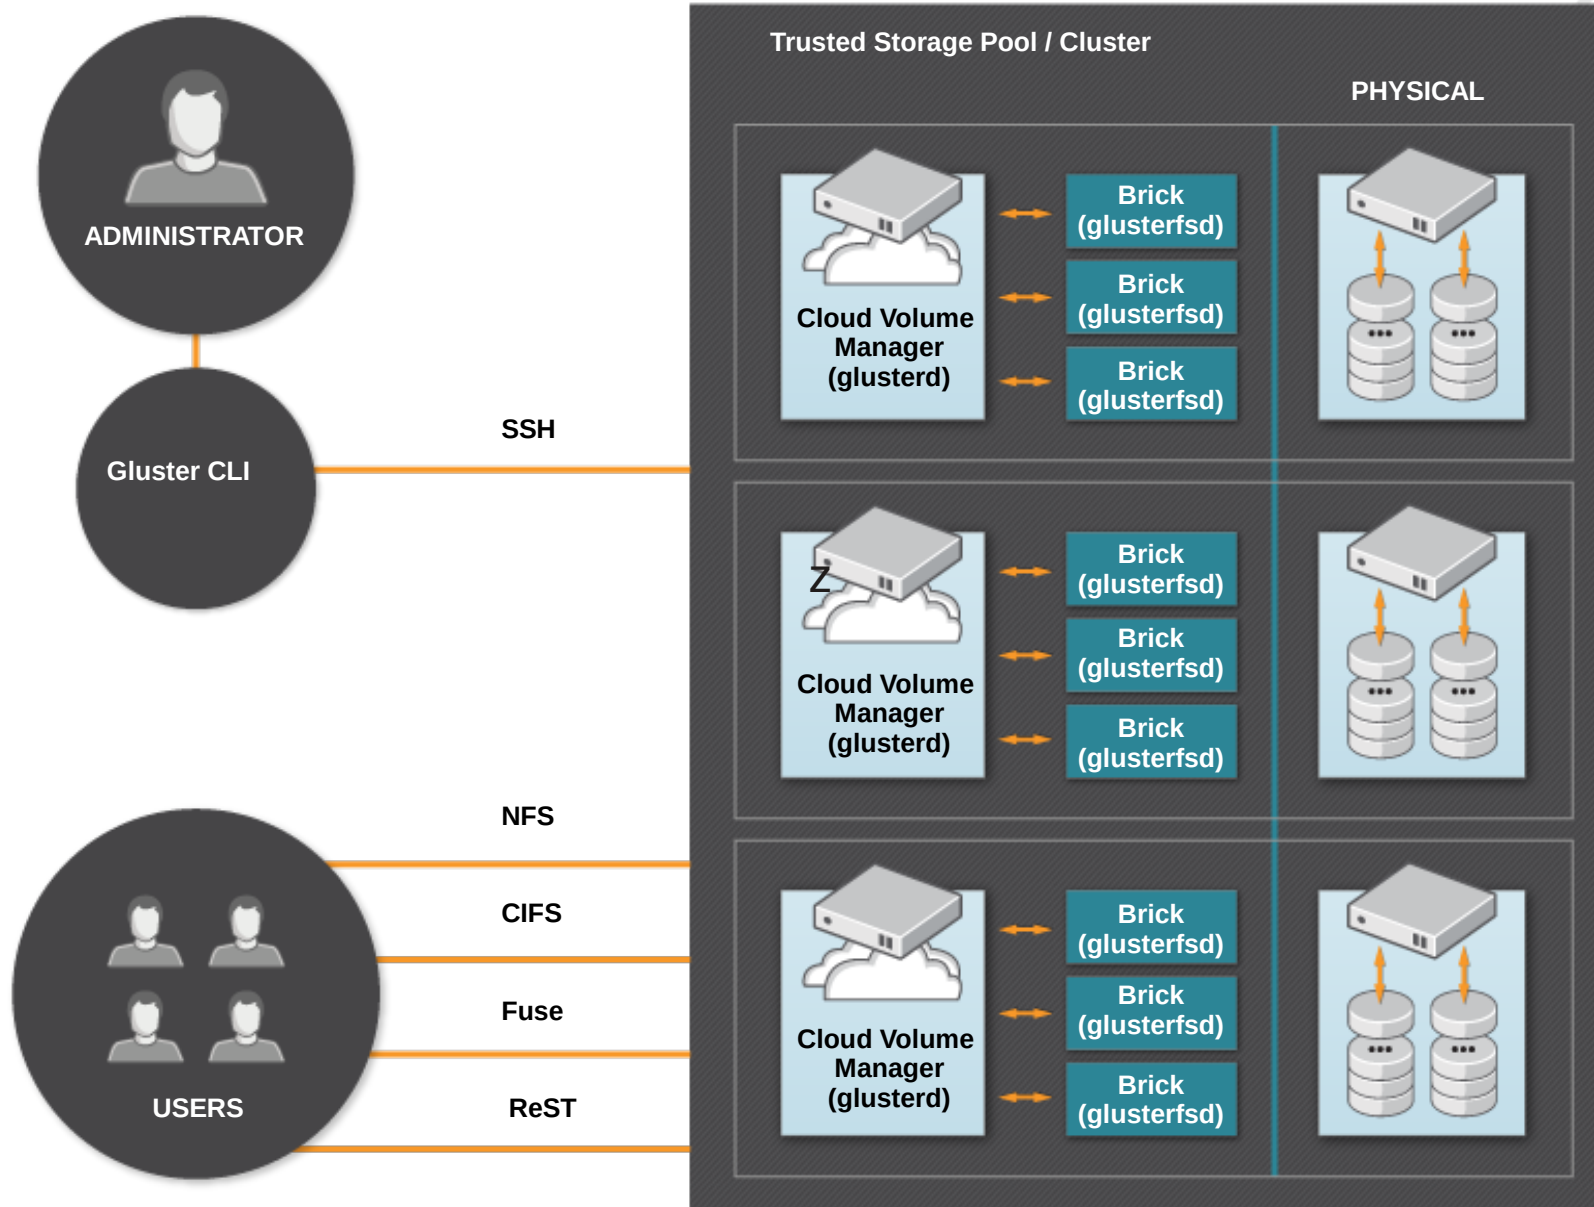
is the basic unit of storage, represented by an export directory on a server



SERVER/NODES

contain the bricks

GlusterFS Overview



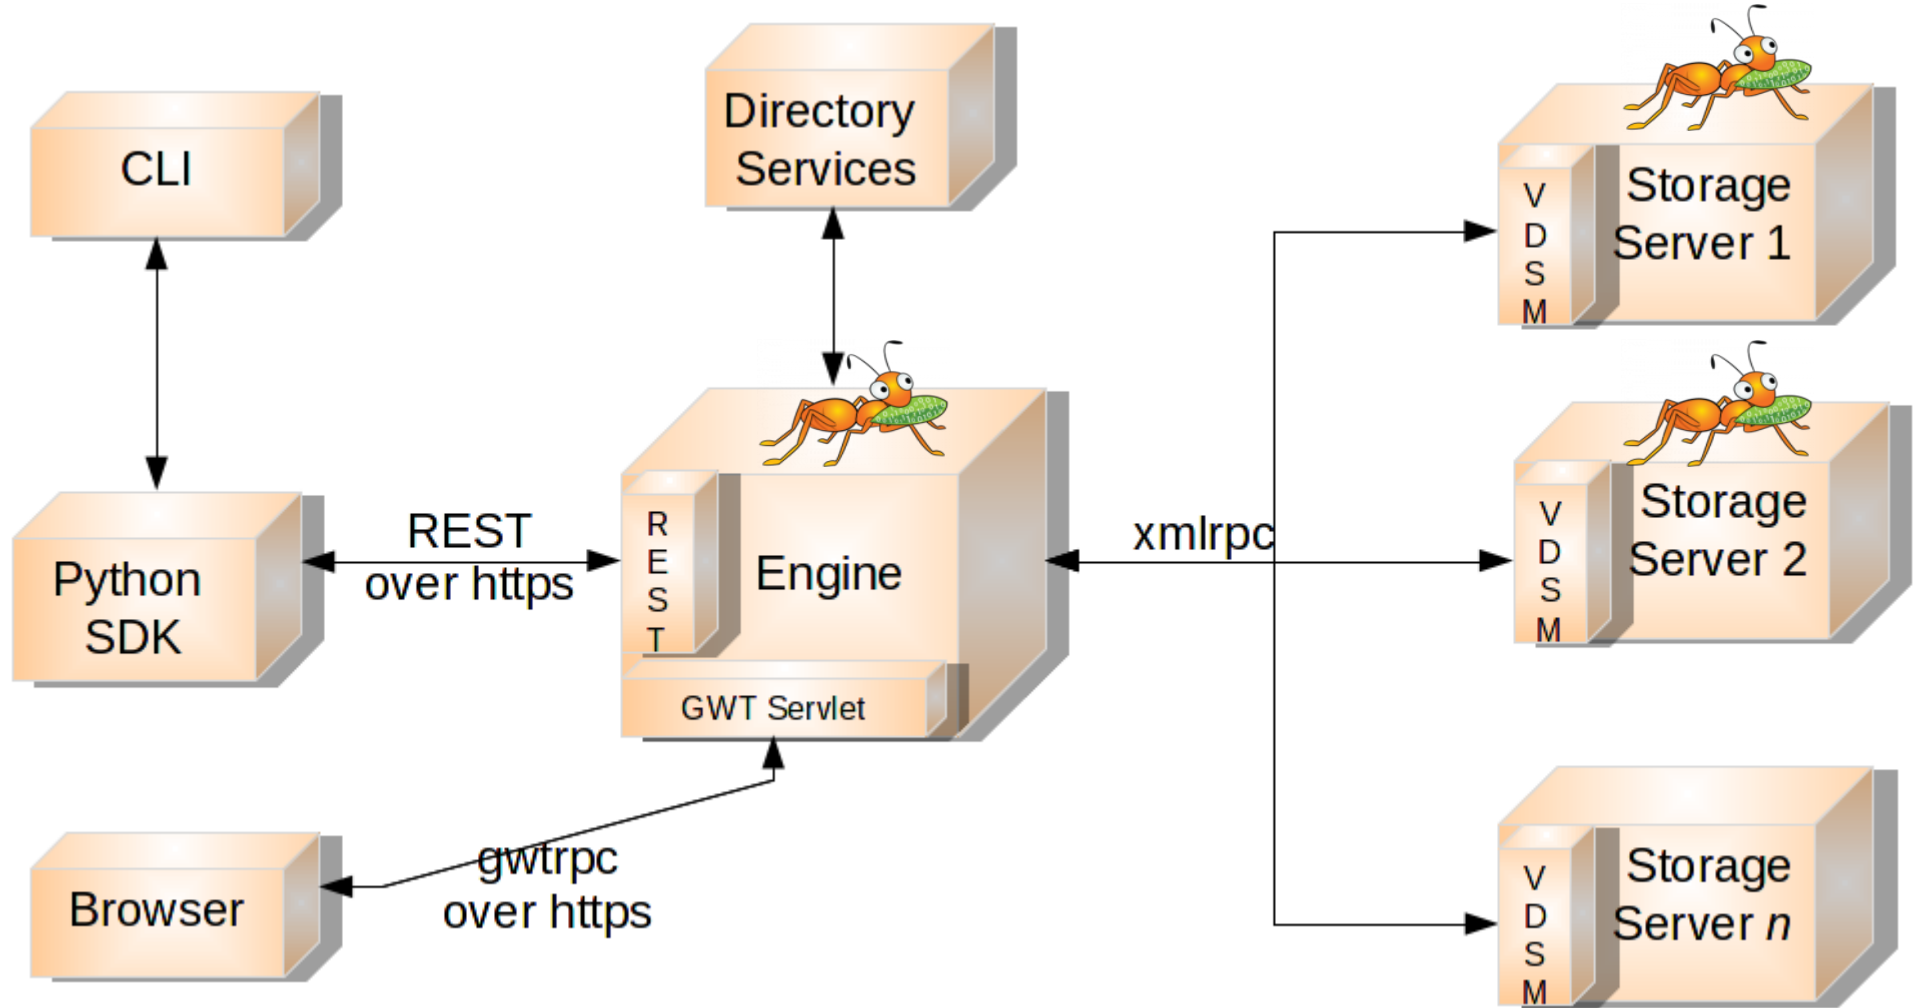
oVirt Gluster Integration



- New feature in oVirt 3.1^[1]
- ApplicationMode configuration
 - 1 → Virtualization only (default)
 - 2 → Gluster only
 - 255 → Virtualization + Gluster
- Enable Gluster at cluster level
- New entities (Volumes, Bricks, Volume Options)
- VDSM verbs for gluster management
 - vdsmd-gluster plug-in

^[1]http://wiki.ovirt.org/wiki/Features/Gluster_Support

Architecture



oVirt, Gluster-ized!




- Cluster Management
 - Create Gluster Cluster
 - Add / Remove Storage Servers
 - Delete Cluster
- Volume Management
 - Create Volume
 - Add / Remove bricks
 - Start / Stop volume
 - Delete Volume



oVirt, Gluster-ized!



- Advanced Search
- Multi Level Administration
- Audit & Alerts
- REST api
- Python SDK & CLI





How does it happen?

Gluster Storage Domain

General flow (fits all application modes)

1. Create Data Center (POSIX compliant)
2. Create Cluster
3. Add host (Hypervisor)
4. Create Volume
5. Add bricks
6. Start volume and optimize for virt
7. Create storage domain (POSIX compliant, VFS type glusterFS)



Gluster Storage Domain

- General flow
 - Create Data Center (POSIX compliant)
 - Create Cluster



New Cluster

General

Name: data

Description: Data Cluster

- Add host (Hypervisor / Storage server)
 - make sure your host has
 - Glusterfs-fuse-3.3.0.5+
 - Glusterfs-3.3.0.5+

New Host

General

Host Cluster: data

Name: server1

Address: 10.16.159.15

Root Password:

Automatically configure host firewall: ☒

Gluster Storage Domain



- General flow
 - Create Volume

The screenshot shows the oVirt management interface with the 'Volumes' tab selected. A 'Create Volume' dialog is open, showing the following configuration:

- Volume Cluster: data
- Name: data
- Type: Distribute
- Transport Type: ☒ TCP
- Bricks: Add Bricks (0 bricks selected)
- Access Protocols:
 - Gluster: ☒
 - NFS: ☒
 - CIFS: ☒
- Allow Access From: *
- (Comma separated list of IP addresses/hostnames)

Buttons: OK, Cancel

- Add bricks

The 'Add Bricks' dialog shows the 'Volume Type' set to 'Distribute'. It includes a 'Bricks' section with a table of existing bricks and input fields for adding new ones.

Server	Brick Directory
<input type="checkbox"/> 10.16.159.228	/export/data
<input type="checkbox"/> 10.16.159.18	/export/data

Buttons: Move Up, Move Down




Gluster Storage Domain

- Optimize volume for virt (Sets optimization options on volume)



- Or even: # gluster volume set VOLNAME group virt



Gluster Storage Domain



- Create PosixFS storage domain (VFS type = glusterfs)

The screenshot displays the oVirt Web Administration interface. A 'New Domain' dialog box is open, showing the following configuration:

- Name: data
- Data Center: Default (POSIX compliant FS)
- Domain Function / Storage Type: Data / POSIX compliant FS
- Format: V3
- Use Host: server2
- Path: 10.16.159.228:/data (with a note: Path to device to mount / remote export)
- VFS Type: glusterfs
- Mount Options: (empty field)

The background interface shows the 'Data Centers' tab, a search bar, and a tree view on the left. The bottom status bar indicates a message: 'Host server4 is non-responsive.' and shows 1 alert.



So what's next?

Upcoming Gluster features (3.2 ++)



- Import existing clusters
- Geo-replication
- Volume advanced details
- Gluster lifecycle hooks management
- Async Tasks
- Unified File and Object storage (UFO) for Swift (OpenStack)
- Storage Device management
- Reports
- And more...

Contribute



- Git repository
 - `git://gerrit.ovirt.org/ovirt-engine`
 - `git://gerrit.ovirt.org/vdsm`
- IRC Channel
 - #ovirt on oftc
 - #gluster on freenode



and now is a good time for....
Questions?



THANK YOU !

<http://www.ovirt.org>

<http://www.gluster.org>